REMARKS

Claims 1-4 and 6-37 are currently pending.

Certain claims have been amended to clarify the scope of Applicant's invention. Claim 5 has been canceled. Claims 32-37 have been added. Claims 1, 7, 12, 17, 19, 20, 23, 24, and 31 have been amended to clearly recite that whereas a first controller may be reset, any other controllers are not reset, and to further clarify the claimed subject matter. The amendment to claims 1, 20, 23, 24, and 31 is supported by original claims 1 and 5, and page 6, lines 5-7, and on page 7, lines 18-20, of Applicant's specification. The amendment to claim 7 is supported by the earlier version of claim 7. The amendment to claims 12 and 19 is supported by page 13, lines 17-20, of Applicant's specification. The amendment to claim 17 is supported by an earlier version of claim 17. The amendment to claim 23 is further supported by page 6, lines 15-17, of Applicant's specification. Claim 25 has been amended to write out the well known meaning of the acronyms. The support for new claims 32-37 is found in Figure 3 and pages 10-13 of Applicant's specification. It is respectfully submitted that no new matter has been added.

The present invention solves prior art problems with multiple controllers in which "either all the controllers reset and there is a loss of access to storage devices during the simultaneous reset of the controllers, or only the defective controller resets and there is insufficient information to solve the defect" (page 2, lines 25-28, of Applicant's specification).

The Patent Office rejected claims 1-4, 6-10, 13-15, 20, 21, 23, 24, and 29 under 35 U.S.C. 102(a) as being anticipated by the "Background of the Invention" (BOI).

For a claim to be anticipated by a reference, each and every element of the claim must be disclosed by that reference (MPEP 2131) unless the element is inherent.

Claim 1 recites "A computer program product stored on a computer readable storage medium for maintaining data access during failure of a first controller in a multiple controller storage subsystem, the storage subsystem having an array of data storage devices and at least one other controller for managing the data storage, comprising computer readable program code for performing the first controller instructing the at least one other controller to save the at least one other controller's internal state information; saving internal state information by the first controller; the first controller resetting itself after the saving of its internal state information; pausing operation of the at least one other controller; the at least one other controller saving

internal state information at the time of pausing, in parallel with the first controller's saving of its internal state information; and continuing operation of the at least one other controller, wherein only the first controller resets."

Claim 20 recites "A method for maintaining data access during failure of a first controller in a multiple controller storage subsystem, the storage subsystem having an array of data storage devices and at least one other controller for managing the data storage, the method comprising the first controller saving internal state information; pausing operation of the at least one other controller; the at least one other controller saving internal state information at the time of pausing without resetting; and, continuing operation of the at least one other controller, wherein only the first controller resets."

Claim 23 recites "An apparatus for maintaining data access during failure of a first controller in a multiple controller storage subsystem, the storage subsystem comprising at least one other controller for managing the data storage, the apparatus comprising the first controller comprising a buffer for storing internal state information when the first controller detects a sanity error in an interface chip; and, the at least one other controller comprising means, responsive to failure of the first controller, for pausing operation, saving internal state information at the time of pausing without resetting, saving information relating to a trace area of the interface chip, and continuing operation, wherein only the first controller resets during an operation for saving the internal state information of the first controller, the internal state information of the at least one other controller, and the information relating to a trace area of the interface chip, wherein interrupts are disabled."

Claim 24 recites "A storage subsystem comprising at least two controllers for managing data storage, the at least two controllers coupled to at least one data storage device, the storage subsystem further comprising a first controller of the at least two controllers adapted for saving internal state information during a failure of the first controller; and, at least one other controller of the at least two controllers adapted for pausing operation, saving internal state information at the time of pausing, and continuing operation during the failure of the first controller, wherein only the first controller resets."

As noted in the BOI, a problem with the prior art, including that discussed in the BOI, is that the problem is "either all the controllers reset and there is a loss of access to storage devices

during the simultaneous reset of the controllers, or only the defective controller resets and there is insufficient information to solve the defect." It is during such failures, the BOI teaches that access to the storage devices is prevented or insufficient information is saved so as to solve the defect or error. The BOI neither discloses nor suggests that a first controller's internal state information is saved and at least one other controller's internal state information is saved without resetting the at least one other controller.

Claim 5 is also now pertinent to these rejected claims because its limitations have been incorporated into the base claims. The Patent Office rejected claim 5 under 35 U.S.C. 103(a) as being unpatentable over BOI in view of Okazaki, U.S. Patent No. 6,345,332.

The Patent Office asserted (page 7-8, of the Office Action mailed October 11, 2005) "Okazaki, on the other hand, teaches about locating the faulty location and resetting only that faulty location without resetting the dual system as a whole (e.g., see Col. 13, lines 60-65). Accordingly, it would have been obvious to one of ordinary skill in the art at the time of the current invention was made to implement the teaching of Okazaki in the computer program product taught by BOI. In doing so, faults are corrected (i.e., on a faulty controller) without interrupting services provided by the dual system (i.e., the other controller(s)). Therefore, it is being advantageous."

The BOI discloses "storage controllers" (e.g., page 1, lines 8, 25-27, 29); Okazaki discloses "bus controllers" (e.g., col. 1, lines 21-23, 46-51, 54-60). The BOI is focused on a fault within a storage controller (e.g., page 2, lines 25-28); Okazaki discloses a fault that may be in an acting system bus, an I/O device, a cross bus, a standby system bus, or a standby system (Fig. 23). Although Okazaki does not disclose which device performs the fault identification method, it seems likely that this is done in the processor (e.g., Fig. 6A). Okazaki discloses that the register for fault information is used to locate a fault and carry out a reset operation at the fault location without resetting the dual system as a whole. Applicant's storage controllers are recited as saving their own internal state information in which the failed storage controller is reset but other storage controllers do not reset themselves. Okazaki is silent regarding any device, such as an I/O device, that stores internal state information. The BOI discloses (page 2, lines 26-27) the prior art in which "only the defective controller resets and there is insufficient information to solve the defect." Modifying the BOI by Okazaki only leads to the prior art technique disclosed

by the BOI. Neither the BOI nor Okazaki, alone or in combination, disclose or fairly suggest the first and other storage controllers storing their internal state information where only the first storage controller resets. Thus, Claims 1, 20, 23, and 24 are allowable over the prior art of record.

Because claims 1, 20, 23, and 24 are allowable, dependent claims 2-19, 21, 22, and 25-30 are also allowable.

The Patent Office rejected claims 11 and 28 under 35 U.S.C. 103(a) as being unpatentable over BOI. Claims 11 and 28 are at least allowable because they depend from allowable base claims 1 and 24, respectively. Applicant challenges the taking of Office Notice and requests an appropriate teaching (i.e., a reference) and motivation (e.g., preferably from that reference) be provided for these limitations. The prior art of record appears not to provide a teaching or a fair suggest for a single circuit card; thus, claims 11 and 28 are allowable for this additional reason.

The Patent Office rejected claims 16, 17, 22, 25, and 30 under 35 U.S.C. 103(a) as being unpatentable over BOI in view of Otterness et al., U.S. Patent No. 6,601,138. Claims 16, 22, 25, and 30 are allowable because they depend from allowable base claims 1 or 24.

Claim 17 recites "A computer program product as claimed in claim 16, wherein during the at least one other controller pausing operation, saving internal state information at the time of pausing, and continuing operation, interrupts are disabled."

BOI discloses "The feature where a failing controller sends a stop message to other controllers is often disabled in the field because the systems are high availability systems" (page 2, lines 16-18, of Applicant's specification). The feature recited in BOI appears to concern an initial default setting in the system regarding whether to allow the propagation of stop messages and not the disabling of interrupts. Contrary to the assertion by the Patent Office on page 8, last four lines, of the Office Action mailed July 14, 2005, BOI does not appear to disclose or even suggest "interrupts are disabled" as recited in claim 17. As the Office Action mailed October 11, 2005, repeats this position, it is noted that a stop message may be received and processed by the other controllers through a polling process, rather than through the use of interrupts. Thus, claim 17 is allowable over the prior art of record for this additional reason.

The Patent Office rejected claim 18 under 35 U.S.C. 103(a) as being unpatentable over

BOI in view of Otterness, further in view of Skazinski et al., U.S. Patent No. 6,574,790.

Claim 18 recites "A computer program product as claimed in claim 16, wherein a flag is set in a host bus adapter during the saving of internal state information to prevent overlapping saves of internal state information in that adapter."

The Patent Office asserted, in paragraph 5, of the Office Action mailed October 11, 2005, (pages 9-10) "As per claim 18, the combination of BOI and Otterness teaches the claimed invention as described above. However, none of them clearly teach about setting a flag to prevent overlapping saves of internal state information in that adapter. Skazinski teaches that using alternate flag (see line 8, Table 6) which is set to equal to true ("1"), to indicate that an alternate mirror entry 6000 is being used to perform the present mirror cache operation to prevent the problems with respect to mirror operation overlap (e.g., see Col. 22, lines 40-47). Accordingly, it would have been obvious to one of ordinary skill in the art at the time of the current invention was made to implement Skazinski's step of setting the flag in the system taught by BOI and Otterness to avoid overlapping saves of internal state information in that adapter."

Skazinski does not disclose setting a flag to prevent overlapping saves of internal state data in that adapter. Skazinski also appears to be concerned with host write data (e.g., column 14, lines 35-48) and not state information, as claimed. Neither Otterness nor the BOI seem to disclose the use of flags. Thus, claim 18 is allowable over the prior art of record for this additional reason.

The Patent Office rejected claim 19 under 35 U.S.C. 103(a) as being unpatentable over BOI in view of Otterness, further in view of Mason, Jr., et al., U.S. Published Patent Application No. 2003/0135674.

Mason discloses, in paragraph 0064, "FIG. 6 is a diagram illustrating how a Peripheral Component Interconnect (PCI) embodiment is connected between the device I/O interface circuitry on a host bus adapter 8 and a disk storage unit 3 utilizing a host I/O interface 7 cable at the front end of the storage management platform and a device I/O interface 2 cable at the back end. Note that the data flow to and from the host bus adapter 8 is over the I/O bus through the I/O cable 7. No data is transferred over the PCI bus 21. This embodiment, taken together with the PCI embodiment illustrated in FIG. 3 illustrate how each and every embodiment of the subject invention plugs into the I/O bus and not the system bus, regardless of whether the host end of the

I/O bus is implemented as an I/O interface chip on a motherboard or an I/O interface ship on a host bus adapter that plugs into the system bus." Mason does not disclose or suggest that the host bus adapter saves information relating to an interface chip. Thus, claim 19 is allowable over the prior art of record for this additional reason.

The Patent Office rejected claims 12, 26, and 27 under 35 U.S.C. 103(a) as being unpatentable over BOI in view of Oldfield et al., U.S. Published Patent Applications No. 2002/0133743.

Claims 26 and 27 are allowable because it depends from allowable claim 24.

Claim 12 recites "A computer program product as claimed in claim 1, wherein in addition to the internal state information, at least one of the first controller and the at least one other controller save external memory data corresponding to an interface chip trace area." The limitation of claim 12 is not met by either the BOI or Oldfield. Thus, claim 12 is allowable over the prior art of record for this additional reason.

The Patent Office rejected claim 31 under 35 U.S.C. 102(a) as being anticipated by DeKoning et al., U.S. Patent No. 5,933,824 hereinafter, DeKoning in view of BOI, further in view of Okazaki.

Claim 31 recites "A Fibre Channel Arbitrated Loop storage system comprising a first set of disk drives connected to a first set of loops, and a second set of disk drives redundant with the first set of disk drives and connected to a second set of loops; wherein a first adapter is connected to the first set of loops and a second adapter is connected to the second set of loops; each adapter being adapted for issuing a command to the other adapter to save internal status data and not reset itself, wherein each adapter is adapted for saving internal status data and resetting."

For a claim to be anticipated by a reference, each and every element of the claim must be disclosed by that reference (MPEP 2131) unless the element is inherent.

Because it is not clear if the Patent Office meant to reject claim 31 under 35 U.S.C. 102(a) or 35 U.S.C. 103(a) and the Patent Office acknowledged a difference not taught by DeKoning in claim 31, Applicant henceforth will treat the rejection of claim 31 as a 103(a) rejection over DeKoning in view of BOI and further in view of Okazaki.

None of DeKoning, the BOI, and Okazaki disclose or suggest "each adapter being adapted for issuing a command to the other adapter to save internal status data without resetting

and each adapter adapted for saving internal status data and resetting." In fact, BOI teaches away from this limitation since BOI (the prior art) has been identified as having a problem in that state data is lost on a controller's failover or insufficient state data is saved in a system when a controller fails to resolve a defect. DeKoning does not acknowledge a problem with error recovery or data storage access during a controller failover and so is not properly modifiable by the BOI. Although Okazaki does not disclose which device performs the fault identification method, it seems likely that this is done in the processor (e.g., Fig. 6A). Okazaki discloses that the register for fault information is used to locate a fault and carry out a reset operation at the fault location without resetting the dual system as a whole. Applicant's storage controllers are recited as saving their own internal state information in which the failed storage controller is reset but other storage controllers do not reset themselves. Okazaki is silent regarding any device, such as an I/O device, that stores internal state information. The BOI discloses (page 2, lines 26-27) the prior art in which "only the defective controller resets and there is insufficient information to solve the defect." Modifying the BOI by Okazaki only leads to the prior art technique disclosed by the BOI. Neither the BOI nor Okazaki, alone or in combination, disclose or fairly suggest the first and other storage controllers storing their internal state information where only the first storage controller resets. Thus, claim 31 is allowable over the prior art of record.

Regarding the remarks section of the Office Action mailed on October 11, 2005. The Patent Office asserted in Paragraph 9 "As to the remark, Applicant asserted: (a) The BOI neither discloses nor suggests that a first controller's internal state information is saved and at least one other controller's internal state information is saved without resetting the at least one other controller. (b) BOI does not appear to disclose or even suggest "interrupts are disabled" as recited in claim 17. (c) None of BOI and Otterness clearly teach about setting a flag to prevent overlapping saves of internal state information in that adapter. (d) Skazinski does not disclose setting a flag to prevent overlapping saves of internal state data in that adapter. Skazinski also appears to be concerned with host write data (e.g., column 14, lines 35-48) and not state information, as claimed. (e) Neither Otterness nor the BOI seem to disclose the use of flags. (f) Mason does not disclose or suggest that the host bus adapter saves information relating to an interface chip. (g) BOI does not disclose or suggest an external memory and does not suggest or

disclose a need for an external memory. (h) The attempt to modify the BOI by Oldfield is impermissible hindsight reconstruction."

The present invention solves prior art problems with multiple controllers in which "either all the controllers reset and there is a loss of access to storage devices during the simultaneous reset of the controllers, or only the defective controller resets and there is insufficient information to solve the defect" (page 2, lines 25-28, of Applicant's specification). Applicant believes that the claims presented in the last response sufficiently drafted around the prior art recognized by Applicant in the BOI. To facilitate prosecution, Applicant has amended the independent claims to clarify Applicant's invention. Applicant believes all independent claims (i.e., 1, 20, 23, 24, and 31) clearly describe applicant's invention such that the claimed invention cannot be construed as falling into one of the two prior art techniques applicant has recognized in the BOI.

Regarding remark (b), page 2, lines 13-19, in pertinent part of the BOI, recites "The feature where a failing controller sends a stop message to other controllers is often disabled in the field because the systems are high availability systems." Claim 17 recites "A computer program product as claimed in claim 16, wherein during the at least one other controller pausing operation, saving internal state information at the time of pausing, and continuing operation, interrupts are disabled." Contrary to the assertion by the Patent Office on page 8, last four lines, of the Office Action mailed July 14, 2005, BOI does not appear to disclose or even suggest "during the at least one other controller pausing operation, saving internal state information at the time of pausing, and continuing operation, interrupts are disabled" as recited in claim 17. As the Office Action mailed October 11, 2005, repeats this position, it is noted that a stop message may be received and processed by the other controllers through a polling process, rather than through the use of interrupts.

Regarding remarks (c) – (e) (all from paragraph 5 of the Office Action mailed October 11, 2005), the Patent Office asserted that Skazinski (Table 6, line 8, and col. 22, lines 40-47) disclose the limitations of Claim 18. Claim 18 recites "wherein a flag is set in a host bus adapter during the saving of internal state information to prevent overlapping saves of internal state information in that adapter." Skazinski discloses (col. 1, line 64, through col. 2, line 23) "in the event of a single controller failure, the surviving controller is able to take over the tasks that were being performed by the failed controller, and perform those tasks that were scheduled to be performed

by the failed controller. To take over the tasks of a failed controller, a surviving controller must keep track of both the tasks that its partner controller is working on, and the tasks that its partner controller is scheduled to work on before the failure occurs. To illustrate this, consider, for example, that a controller fails before data stored in its cache (in response to a write request from a host computer) is written onto a system drive. Data in the cache of a failed controller is lost unless a battery backup is used. In this situation, it is desirable for a surviving controller to complete the scheduled task of the failed controller by writing the data that was in the failed controller's cache onto the system drive. To accomplish this, a surviving controller in active configuration would need to have a copy, or a mirror of the failed controller's cache. State-of-the-art data storage systems are limited because there are no known structure or procedures for copying or mirroring a controller's cache between other different controllers in active configuration. Therefore, what is needed, is a cache mirroring system, apparatus, and method for multi-controller environments." Skazinski discloses (col. 22, lines 40-47) "Finally, using alternate flag (see line 8, Table 6) is set to equal to true ("1"), to indicate that an alternate mirror entry 6000 is being used to perform the present mirror cache operation to prevent the problems discussed in greater detail above with respect to mirror operation overlap. In this manner, CDMP 300 sets up the mirror operation to mirror to an alternate mirror cache line." Even though Skazinski discloses flags to prevent mirror operation overlap, Skazinski is directed to the mirroring of cache data of an adapter so that the tasks of any controller are not lost and does not disclose the setting of a flag within an adapter "wherein a flag is set in a host bus adapter during the saving of internal state information to prevent overlapping saves of internal state information in that adapter." That is, even though Skazinski discloses flags, the flags of Skazinski prevent mirror data overlaps and are not set within an adapter to prevent overlapping saves of internal state information in that adapter. Thus, claim 18 is not made obvious by the BOI, Otterness, and Skazinski, alone or in combination.

Regarding remark (f), claim 19 recites "the host bus adapter saves information relating to an interface chip trace area." Even though Mason discloses an interface, Mason does not seem to disclose an interface chip, an interface chip trace area, or the host bus adapter saves information relating to an interface chip trace area. Even if Mason (paragraph 0064) is construed so as to make obvious an interface chip, there appears to be no disclosure and no fair suggestion by

Otterness, the BOI, or Mason "that the host bus adapter saves information relating to an interface chip trace area." Thus, claim 19 is allowable for this additional reason.

Regarding remarks (g) and (h), claim 12 is allowable on its own merits over the BOI and Oldfield as discussed above in the rejection. Claims 26 and 27 are allowable because they depend from an allowable base claim.

The Patent Office is respectfully requested to reconsider and remove the rejections of the claims under 35 U.S.C. 102(a) or 103(a) based on BOI, Otterness, DeKoning, Oldfield, Skazinski, Mason, and Okazaki, alone or in combination, and to allow all of the pending claims 1-4 and 6-37 as now presented for examination. An early notification of the allowability of claims 1-4 and 6-37 is earnestly solicited.

S.N.: 10/608,721
Art Unit: 2186

Respectfully submitted:

Walter J. Malinowski December 13, 2005

Walter J. Malinowski Date

Reg. No.: 43,423

Customer No.: 29683

HARRINGTON & SMITH, LLP
4 Research Drive
Shelton, CT 06484-6212

Telephone: (203)925-9400
Facsimile: (203)944-0245
email: wmalinowski@hspatent.com

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, P.O. BOX 1450, Alexandria, VA 22313-1450.

12/14/2005 Elaine F. Mion

Date Name of Person Making Deposit